

(12) **United States Patent**
Benhase et al.

(10) **Patent No.:** **US 9,471,496 B2**
(45) **Date of Patent:** **Oct. 18, 2016**

(54) **DEMOTING TRACKS FROM A FIRST CACHE TO A SECOND CACHE BY USING A STRIDE NUMBER ORDERING OF STRIDES IN THE SECOND CACHE TO CONSOLIDATE STRIDES IN THE SECOND CACHE**

USPC 711/118, 119, 122
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **International Business Machines Corporation**, Armonk, NY (US)

4,464,713 A 8/1984 Benhase et al.
5,860,090 A 1/1999 Clark

(Continued)

(72) Inventors: **Michael T. Benhase**, Tucson, AZ (US);
Lokesh M. Gupta, Tucson, AZ (US)

FOREIGN PATENT DOCUMENTS

(73) Assignee: **INTERNATIONAL BUSINESS MACHINES CORPORATION**,
Armonk, NY (US)

CN 1967495 5/2007
CN 1967507 5/2007
WO 2011042428 4/2011

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 116 days.

OTHER PUBLICATIONS

(21) Appl. No.: **14/465,769**

Response dated Jul. 14, 2014, pp. 13, to Final Office Action dated Apr. 7, 2014, pp. 41, U.S. Appl. No. 13/352,230.

(22) Filed: **Aug. 21, 2014**

(Continued)

(65) **Prior Publication Data**

US 2014/0365718 A1 Dec. 11, 2014

Related U.S. Application Data

(63) Continuation of application No. 13/779,412, filed on Feb. 27, 2013, now Pat. No. 8,825,956, which is a continuation of application No. 13/352,236, filed on Jan. 17, 2012, now Pat. No. 8,825,953.

(51) **Int. Cl.**

G06F 12/00 (2006.01)
G06F 13/00 (2006.01)

(Continued)

(52) **U.S. Cl.**

CPC **G06F 12/0833** (2013.01); **G06F 3/0619** (2013.01); **G06F 3/0667** (2013.01);
(Continued)

(58) **Field of Classification Search**

CPC G06F 12/0246; G06F 12/0833; G06F 12/122

Primary Examiner — Charles Rones

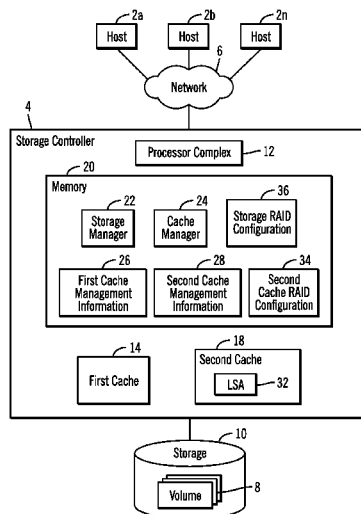
Assistant Examiner — Sidney Li

(74) *Attorney, Agent, or Firm* — David W. Victor; Konrad Raynes Davda & Victor LLP

(57) **ABSTRACT**

Information on strides configured in the second cache includes information indicating a number of valid tracks in the strides, wherein a stride has at least one of valid tracks and free tracks not including valid data. A determination is made of tracks to demote from the first cache. A first stride is formed including the determined tracks to demote. The tracks from the first stride are added to a second stride in the second cache that has no valid tracks. A target stride in the second cache is selected based on a stride most recently used to consolidate strides from at least two strides into one stride. Data from the valid tracks is copied from at least two source strides in the second cache to the target stride.

22 Claims, 10 Drawing Sheets



- (51) **Int. Cl.**
G06F 13/28 (2006.01)
G06F 12/08 (2016.01)
G06F 12/12 (2016.01)
G06F 12/02 (2006.01)
G06F 3/06 (2006.01)
- (52) **U.S. Cl.**
 CPC **G06F3/0689** (2013.01); **G06F 12/0246**
 (2013.01); **G06F 12/122** (2013.01); **G06F**
12/123 (2013.01); **G06F 12/128** (2013.01);
G06F 12/0871 (2013.01); **G06F 12/0897**
 (2013.01); **G06F 2212/225** (2013.01); **G06F**
2212/69 (2013.01); **G06F 2212/7209** (2013.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,860,091	A *	1/1999	DeKoning	G06F 3/0626	711/113
6,321,240	B1	11/2001	Chilimbi et al.			
6,785,771	B2	8/2004	Ash et al.			
7,107,385	B2	9/2006	Rajan et al.			
7,444,478	B2	10/2008	LaFrese et al.			
7,562,203	B2	7/2009	Scott et al.			
7,650,341	B1	1/2010	Oratovsky et al.			
7,669,022	B2	2/2010	Maruyama et al.			
7,698,501	B1	4/2010	Corbett et al.			
7,702,870	B2	4/2010	English et al.			
7,721,043	B2	5/2010	Gill et al.			
7,725,651	B2	5/2010	Saito			
7,769,960	B2	8/2010	LaFrese et al.			
7,930,325	B2	4/2011	Sieglwart et al.			
2001/0029574	A1	10/2001	Razdan			
2002/0166022	A1 *	11/2002	Suzuki	G06F 12/0866	711/103
2002/0199070	A1	12/2002	Chaudhry			
2003/0028695	A1	2/2003	Burns et al.			
2003/0070042	A1	4/2003	Byrd et al.			
2003/0105928	A1 *	6/2003	Ash	G06F 12/0804	711/136
2004/0068612	A1	4/2004	Stolowitz			
2004/0098541	A1	5/2004	Megiddo et al.			
2005/0073884	A1	4/2005	Gonzalez et al.			
2005/0144360	A1	6/2005	Bennett et al.			
2005/0257083	A1 *	11/2005	Cousins	G06F 17/30067	714/6.2
2006/0106891	A1	5/2006	Mahar et al.			
2006/0155934	A1	7/2006	Rajamony et al.			
2007/0106707	A1 *	5/2007	Yamato	G06F 12/0246	
2007/0118695	A1	5/2007	Lowe et al.			
2009/0210620	A1	8/2009	Jibbe et al.			
2009/0216954	A1	8/2009	Benhase et al.			
2009/0271412	A1	10/2009	Lacapra et al.			
2009/0271418	A1	10/2009	Vaghani et al.			
2010/0082774	A1	4/2010	Pitts			
2010/0166022	A1	7/2010	Cho			
2010/0191907	A1 *	7/2010	Ish	G06F 11/1096	711/114
2010/0293420	A1	11/2010	Kapil et al.			
2011/0087837	A1	4/2011	Blinick et al.			
2011/0191523	A1	8/2011	Caulkins			
2011/0202732	A1	8/2011	Montgomery			
2013/0111106	A1	5/2013	Benhase et al.			
2013/0111134	A1	5/2013	Benhase et al.			

OTHER PUBLICATIONS

Notice of Allowance dated Oct. 6, 2014, pp. 21, for U.S. Appl. No. 13/352,230.
 Notice of Allowance dated Sep. 30, 2014, pp. 14, for U.S. Appl. No. 13/464,668.
 Office Action dated Nov. 8, 2013, pp. 65 for U.S. Appl. No. 13/352,230.

Response dated Feb. 7, 2014, pp. 13, to Office Action dated Nov. 8, 2013, pp. 65 for U.S. Appl. No. 13/352,230.
 Final Office Action dated Mar. 24, 2014, pp. 37, for U.S. Appl. No. 13/352,230.
 Office Action dated Nov. 8, 2013, pp. 65 for U.S. Appl. No. 13/464,668.
 Response dated Feb. 7, 2014, pp. 9, to Office Action dated Nov. 8, 2013, pp. 65 for U.S. Appl. No. 13/464,668.
 Final Office Action dated Mar. 24, 2014, pp. 36 for U.S. Appl. No. 13/464,668.
 Office Action dated Aug. 30, 2013, pp. 55, for U.S. Appl. No. 13/352,224, filed Jan. 17, 2012.
 Response dated Dec. 2, 2013, p. 13 to Office Action dated Aug. 30, 2013, pp. 55, for U.S. Appl. No. 13/352,224, filed Jan. 17, 2012, by inventors M. et al.
 Notice of Allowance dated Apr. 29, 2014, pp. 55, for U.S. Appl. No. 13/352,224, filed Jan. 17, 2012.
 Office Action dated Sep. 5, 2013, pp. 33, for U.S. Appl. No. 13/779,309, filed Feb. 27, 2013.
 Response dated Dec. 5, 2013, pp. 9, to Office Action dated Sep. 9, 2013, pp. 33, for U.S. Appl. No. 13/779,309, filed Feb. 27, 2013.
 Notice of Allowance dated Apr. 30, 2014, pp. 50, for U.S. Appl. No. 13/779,309, filed Feb. 27, 2013.
 Office Action dated Oct. 3, 2013, pp. 37, for U.S. Appl. No. 13/352,236, filed Jan. 17, 2012.
 Response dated Jan. 3, 2014, pp. 1, to Office Action dated Oct. 3, 2013, pp. 37, for U.S. Appl. No. 13/352,236, filed Jan. 17, 2012.
 Office Action dated Jan. 10, 2014, pp. 67, for U.S. Appl. No. 13/352,236, filed Jan. 17, 2012.
 Response dated Apr. 8, 2014, pp. 10, to Office Action dated Jan. 10, 2014, pp. 67, for U.S. Appl. No. 13/352,236, filed Jan. 17, 2012.
 Notice of Allowance dated Apr. 30, 2014, pp. 29, for U.S. Appl. No. 13/352,236, filed Jan. 17, 2012.
 U.S. Appl. No. 13/779,412, filed Feb. 27, 2013, by inventors M. Benhase et al.
 Preliminary Remarks filed Feb. 27, 2013, pp. 2, for U.S. Appl. No. 13/779,412, filed Feb. 27, 2013, by inventors M. Benhase et al.
 Office Action dated Oct. 10, 2013, pp. 45, for U.S. Appl. No. 13/779,412, filed Feb. 27, 2013.
 Response dated Jan. 10, 2014, pp. 9, for U.S. Appl. No. 13/779,412, filed Feb. 27, 2013.
 Notice of Allowance dated Apr. 22, 2014, pp. 45, U.S. Appl. No. 13/779,412, filed Feb. 27, 2013.
 Response dated Jul. 14, 2014, pp. 10, to Final Office Action dated Apr. 7, 2014, pp. 41, for U.S. Appl. No. 13/464,668.
 Office Action dated Oct. 18, 2013, pp. 54, for U.S. Appl. No. 13/352,239, filed Jan. 17, 2012.
 Response dated Jan. 21, 2014, pp. 13, to Office Action dated Oct. 18, 2013, pp. 54, for U.S. Appl. No. 13/352,239, filed Jan. 17, 2012.
 Final Office Action dated Apr. 7, 2014, pp. 41, for U.S. Appl. No. 13/352,239, filed Jan. 17, 2012.
 Response dated Jul. 7, 2014, pp. 13, to Final Office Action dated Apr. 7, 2014, pp. 41, for U.S. Appl. No. 13/352,239, filed Jan. 17, 2012.
 Office Action dated Oct. 25, 2013, pp. 49, for U.S. Appl. No. 13/779,439, filed Feb. 27, 2013.
 Response dated Jan. 21, 2014, pp. 10, to Office Action dated Oct. 25, 2013, pp. 49, for U.S. Appl. No. 13/779,439, filed Feb. 27, 2013.
 Final Office Action dated Apr. 4, 2014, pp. 38, for U.S. Appl. No. 13/779,439, filed Feb. 27, 2013.
 Response dated Jul. 7, 2014, pp. 10, to Final Office Action dated Apr. 4, 2014, pp. 38, for U.S. Appl. No. 13/779,439, filed Feb. 27, 2013.
 U.S. Appl. No. 13/352,230, filed Jan. 17, 2012, entitled "Populating a First Stride of Tracks from a First Cache to Write to a Second Stride in a Second Cache", by inventors K. Ash, M. Benhase, L. Gupta, M. Kalos, K. Nielsen.
 U.S. Appl. No. 13/352,224, filed Jan. 17, 2012, entitled "Demoting Tracks from a First Cache to a Second Cache by Using an Occupancy of Valid Tracks in Strides in the Second Cache to Consolidate Strides in the Second Cache", by inventors M. Benhase, and L. Gupta.

(56)

References Cited**OTHER PUBLICATIONS**

U.S. Appl. No. 13/352,239, filed Jan. 17, 2012, entitled "Demoting Partial Tracks From a First Cache to a Second Cache", by inventors K. Ash, M. Benhase, and L. Gupta.

U.S. Appl. No. 13/133,974, filed May 23, 2011, entitled "Cache Management of Tracks in a First Cache and a Second Cache for Storage", by inventors M. Benhase, et al.

U.S. Appl. No. 13/113,931, filed May 23, 2011, entitled "Populating Strides of Tracks to Demote From a First Cache to a Second Cache", by inventors M. Benhase, et al.

U.S. Appl. No. 13/113,937, filed May 23, 2011, entitled "Managing Unmodified Tracks Maintained in Both a First Cache and a Second Cache", by inventors K. Ash, et al.

U.S. Appl. No. 13/113,944, filed May 23, 2011, entitled "Caching Data in a Storage System Having Multiple Caches Including Non-Volatile Storage Cache in a Sequential Access Storage Device", by inventors M. Benhase, et al.

U.S. Appl. No. 13/113,949, filed May 23, 2011, entitled "Using an Attribute of a Write Request to Determine Where to Cache Data in a Storage System Having Multiple Caches Including Non-Volatile Storage Cache in a Sequential Access Storage Device", by inventors M. Benhase, et al.

U.S. Appl. No. 13/113,953, filed May 23, 2011, entitled "Handling High Priority Requests in a Sequential Access Storage Device Having a Non-Volatile Storage Cache", by inventors M. Benhase, et al.

U.S. Appl. No. 13/113,958, filed May 23, 2011, entitled "Handling High Priority Requests in a Sequential Access Storage Device Having a Non-Volatile Storage Cache", by inventors M. Benhase, et al.

"Check Point for a Two Stage Store", IBM Corp., IP.com Document No. IPCOM000089366D, TDB Oct. 1977, pp. 1955-1958, Oct. 1, 1977, pp. 1-5.

"Serial ATA Native Command Queuing" joint White Paper by Intel Corporation and Seagate Technology, Jul. 2003, pp. 1-12.

"Multiple Command Control and Reordering", [online] [retrieved May 14, 2011], pp. 1-2 <http://www.pcguides.com/ref/hdd/op/logicMultiple-c.html>.

"Superparamagnetic Effect on Hard Disk Drive", [online] [retrieved May 16, 2011], pp. 1-2, <http://www.dataclinic.co.uk/hard-disk-superparamagnetic-effect.html>.

"Superparamagnetism", Wikipedia, [online] [retrieved May 19, 2011], pp. 1-5, <http://en.wikipedia.org/w/index.php?title=Superparamagnetism&printable>

"Seagate's Terabyte Platters Make it the Densest of the Lot", The Register, [online] [retrieved May 19, 2011], pp. 1-2, http://www.theregister.co.uk/2011/05/03/seagate_terabyte_platter/

Hitachi Demos 230 Gb Per Square Inch Data Density on Perpendicular Re . . . , [online] [retrieved May 19, 2011], pp. 1-9, <http://www.physorg.com/news3588.html>.

"Hard Disk Drive", Wikipedia, [online] [retrieved May 19, 2011], pp. 1-23, http://en.wikipedia.org/w/index.php?title=hard_disk_drive&printable=yes.

Preliminary Amendment 1, filed Jul. 12, 2011, 16 pp., for U.S. Appl. No. 13/113,931, filed May 23, 2011 by M.T. Benhase et al.

Preliminary Amendment 2, filed May 7, 2012, 8 pp., for U.S. Appl. No. 13/113,931, filed May 23, 2011 by M.T. Benhase et al.

U.S. Appl. No. 13/465,717, filed May 7, 2012, by inventors M. Benhase, et al.

Preliminary Amendment, pp. 20, filed May 7, 2012, for U.S. Appl. No. 13/465,717, by M.T. Benhase et al.

Preliminary Amendment filed May 4, 2012, 8 pp., for U.S. Appl. No. 13/352,230, by inventors K. Ash, et al.

U.S. Appl. No. 13/464,668, filed May 4, 2012, by inventors K. Ash, et al.

Preliminary Amendment filed May 4, 2012, 6 pp., for U.S. Appl. No. 13/464,668, by inventors K. Ash, et al.

U.S. Appl. No. 13/352,224, filed Jan. 17, 2012, by inventors M. et al.

Preliminary Amendment filed Feb. 27, 2013, pp. 9, for U.S. Appl. No. 13/352,224, filed Jan. 17, 2012, by inventors M. et al.

U.S. Appl. No. 13/779,309, filed Feb. 27, 2013, by inventors M. et al.

Preliminary Remarks filed Feb. 27, 2013, pp. 2, for U.S. Appl. No. 13/779,309, filed Feb. 27, 2013, by inventors M. et al.

U.S. Appl. No. 13/352,236, filed Jan. 17, 2012, by inventors M. Benhase et al.

Preliminary Amendment filed Feb. 27, 2013, pp. 8, for U.S. Appl. No. 13/352,236, filed Jan. 17, 2012, by inventors M. Benhase et al.

U.S. Appl. No. 13/352,239, filed Jan. 17, 2012, by inventors K. Ash, et al.

Preliminary Amendment filed Feb. 27, 2013, pp. 8, for U.S. Appl. No. 13/352,239, filed Jan. 17, 2012, by inventors K. Ash, et al.

U.S. Appl. No. 13/779,439, filed Feb. 27, 2013, by inventors K. Ash, et al.

Preliminary Remarks filed Feb. 27, 2013, pp. 2, for U.S. Appl. No. 13/779,439, filed Feb. 27, 2013, by inventors K. Ash, et al.

PCT Search Report & Written Opinion dated May 2, 2013 for PCT/IB2012/057140 filed Dec. 10, 2012.

English Translation of CN1967495 filed May 23, 2007 by IBM Corp.

English Translation of CN1967507 filed May 23, 2007 by IBM Corp.

UK Combined Search Report & Examination Report dated Jul. 19, 2013 for GB1300444.5 filed Jan. 10, 2013.

UK Response dated Dec. 9, 2013 to Official Letter received Jul. 19, 2013 for GB1300444.5 filed Jan. 10, 2013.

Preliminary Amendment 2, filed May 7, 2012, 8 pp., for U.S. Appl. No. 13/113,931, filed on May 23, 2011.

Office Action dated Jul. 8, 2013, pp. 31, for U.S. Appl. No. 13/113,931, filed May 23, 2011.

Response dated Oct. 8, 2013, pp. 12, to Office Action dated Jul. 8, 2013, pp. 31, for U.S. Appl. No. 13/113,931, filed May 23, 2011.

Office Action dated Dec. 24, 2013, pp. 38, for U.S. Appl. No. 13/113,931, filed May 23, 2011 by M.T. Benhase et al.

Response dated Mar. 24, 2014, pp. 10, to Office Action dated Dec. 24, 2013, pp. 38, for U.S. Appl. No. 13/113,931, filed May 23, 2011.

Notice of Allowance dated Apr. 17, 2014, pp. 18, for U.S. Appl. No. 13/113,931, filed May 23, 2011.

Response dated Oct. 8, 2013 to Office Action dated Jul. 9, 2013, pp. 28, for U.S. Appl. No. 13/465,717.

Notice of Allowance dated Jan. 28, 2014, pp. 14, for U.S. Appl. No. 13/465,717.

Office Action dated Jul. 9, 2013, pp. 28, for U.S. Appl. No. 13/465,717.

Notice of Allowance received Dec. 23, 2014, pp. 31, for U.S. Appl. No. 13/352,239, filed Jan. 17, 2012.

Notice of Allowance dated Dec. 26, 2014, pp. 20, for U.S. Appl. No. 13/779,439, filed Feb. 27, 2013.

* cited by examiner

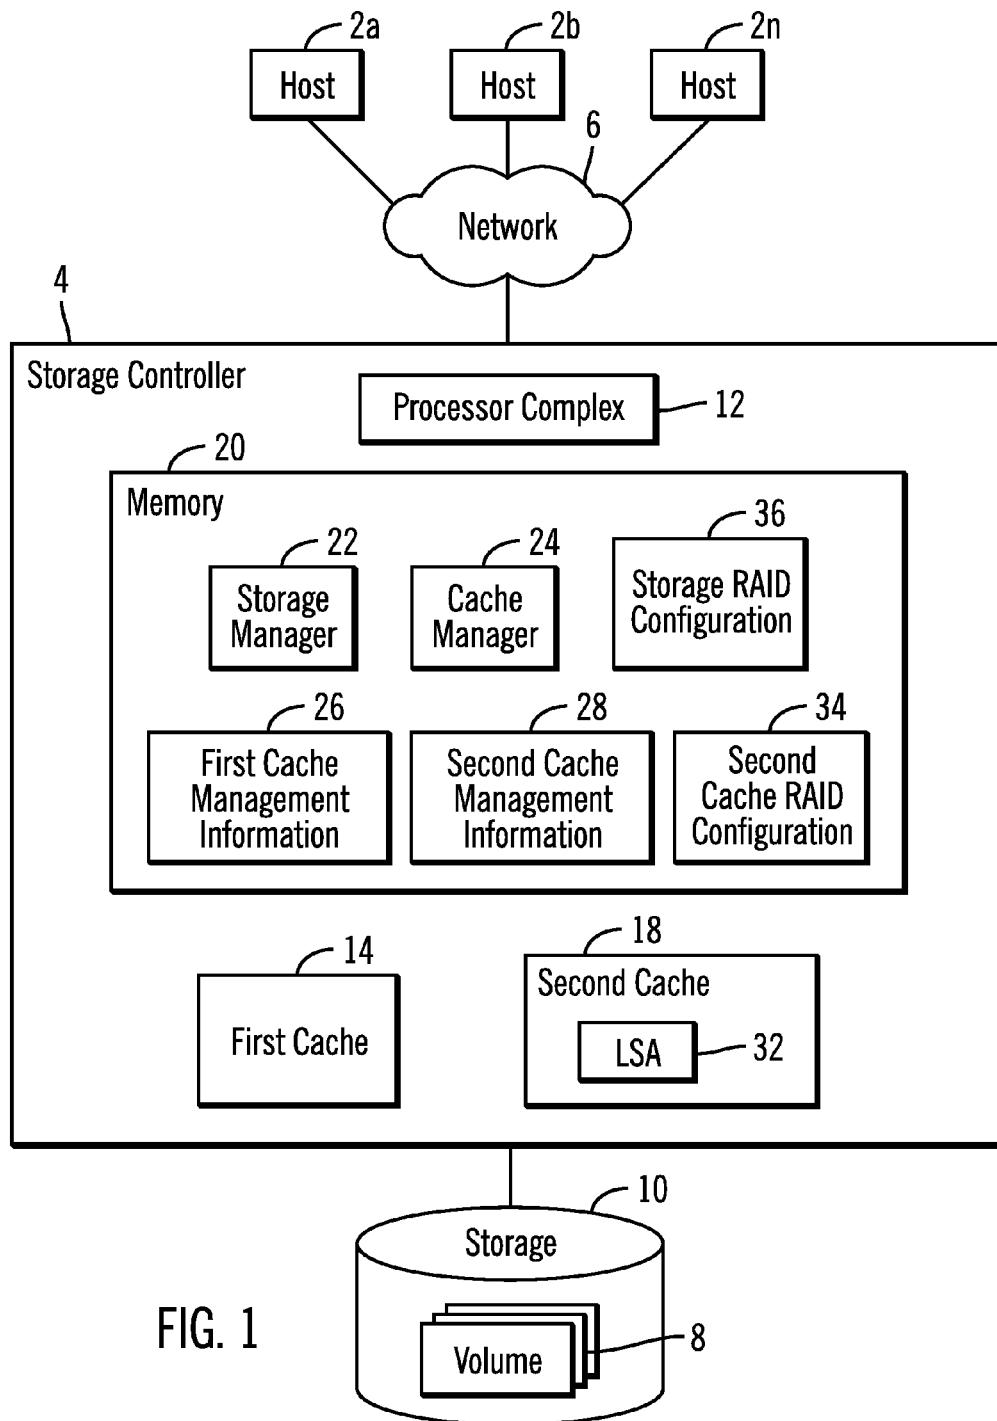


FIG. 1

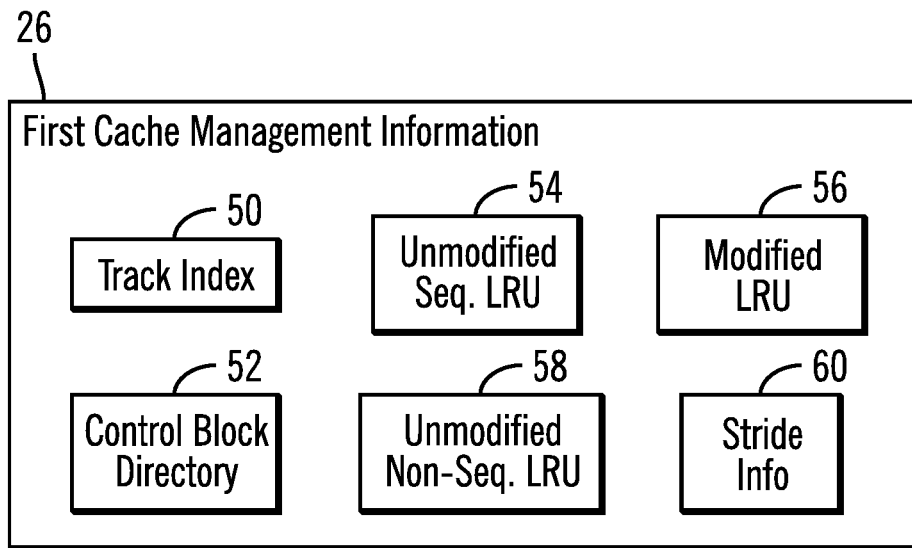


FIG. 2

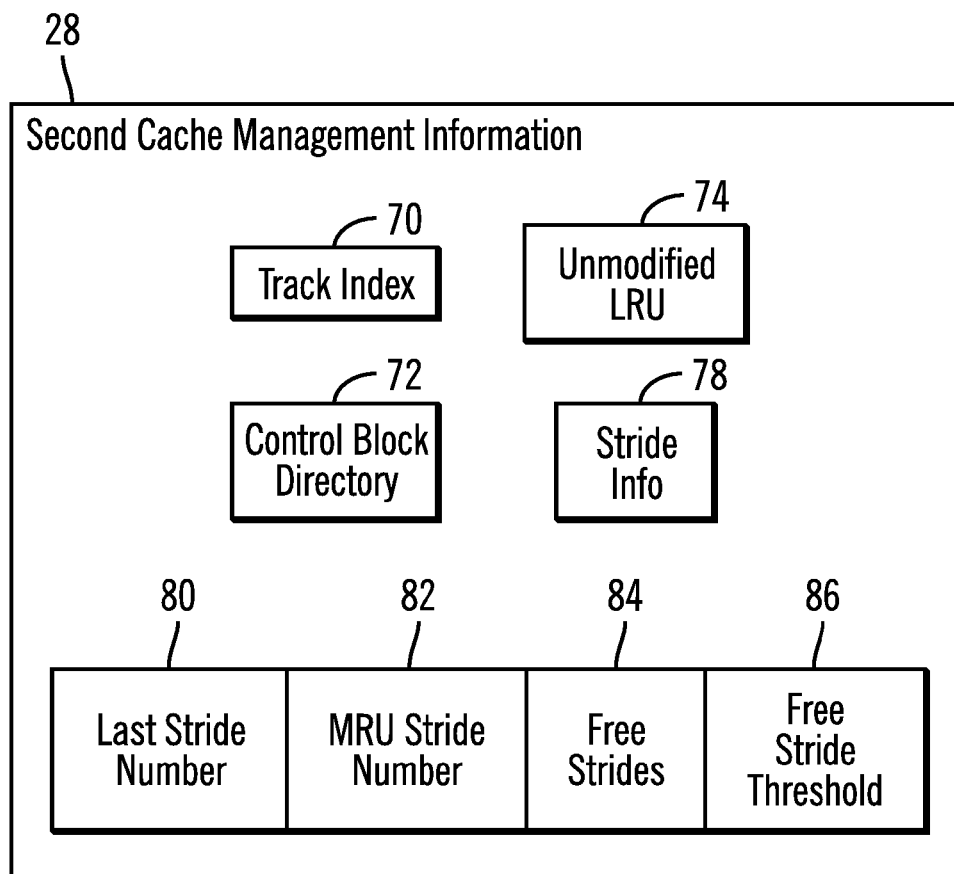
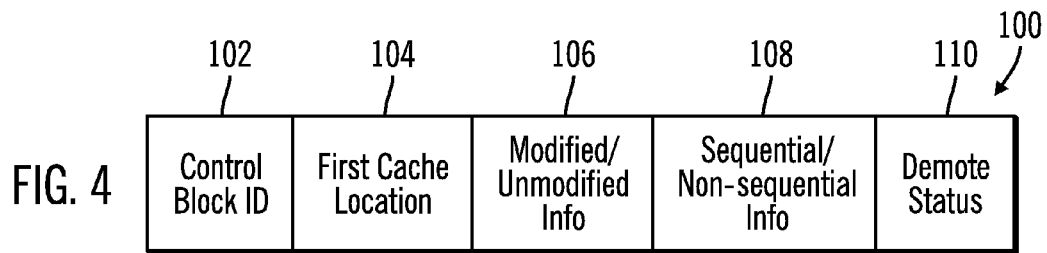
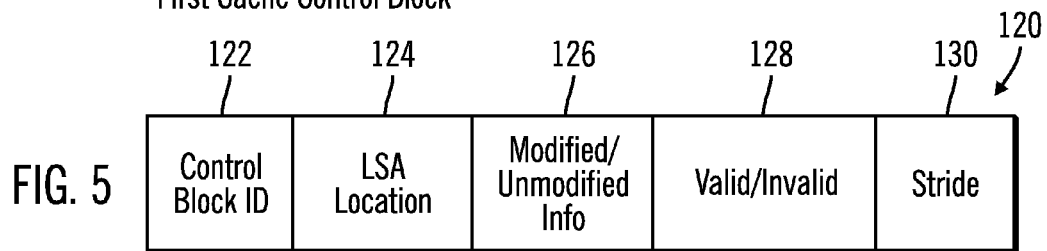


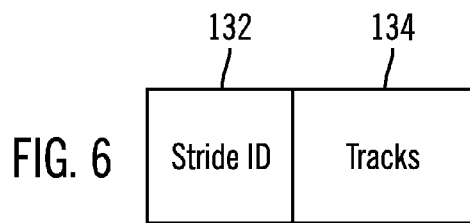
FIG. 3



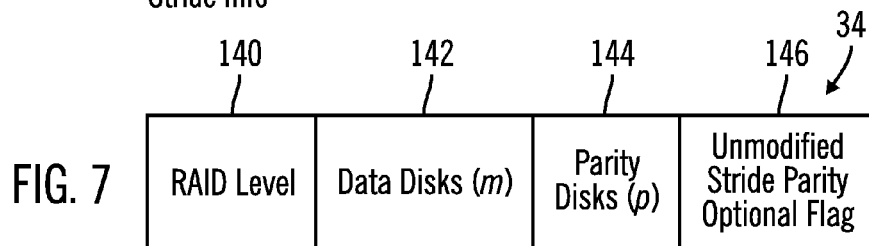
First Cache Control Block



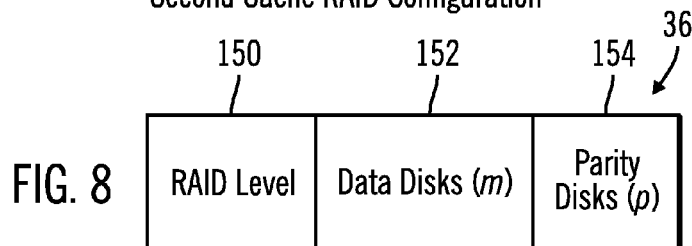
Second Cache Control Block



Stride Info



Second Cache RAID Configuration



Storage System RAID Configuration

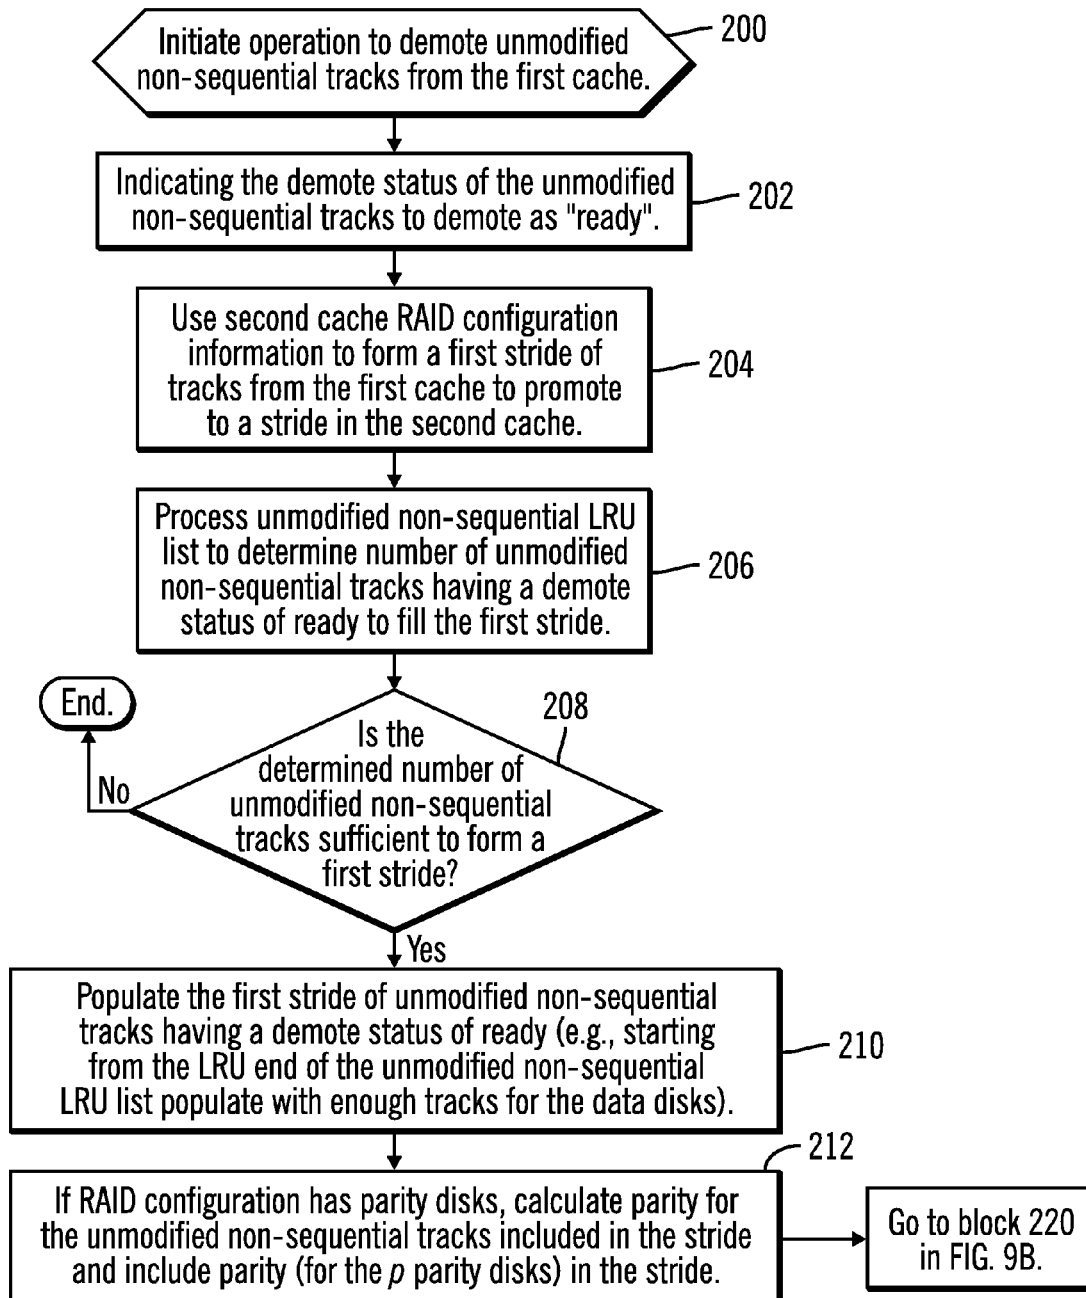


FIG. 9A

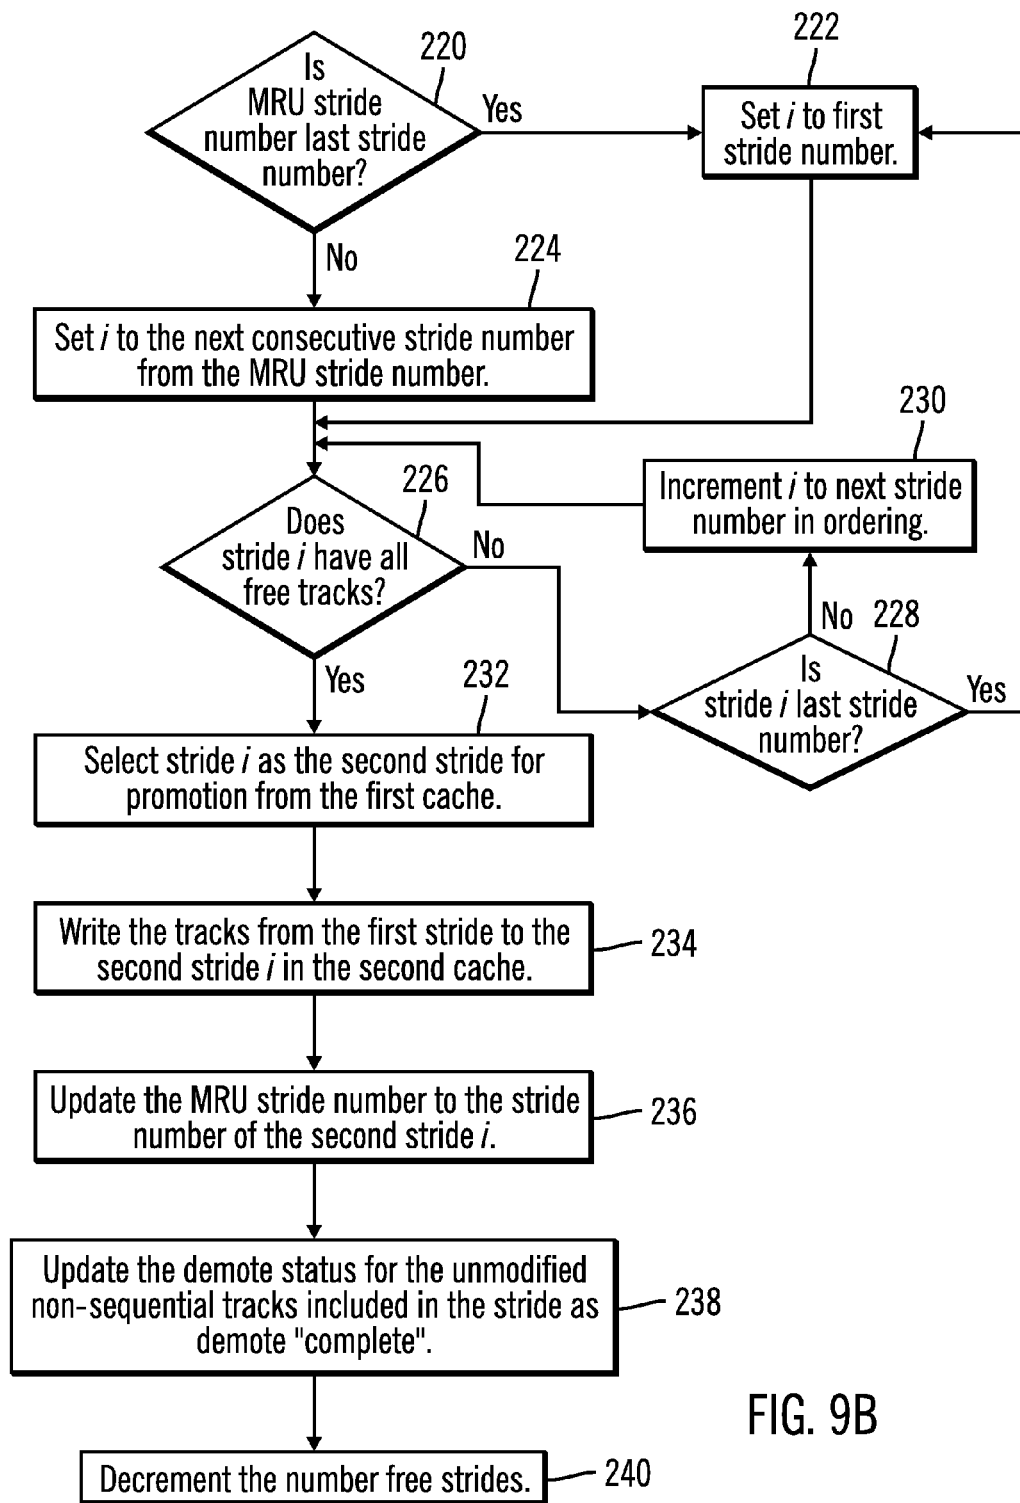


FIG. 9B

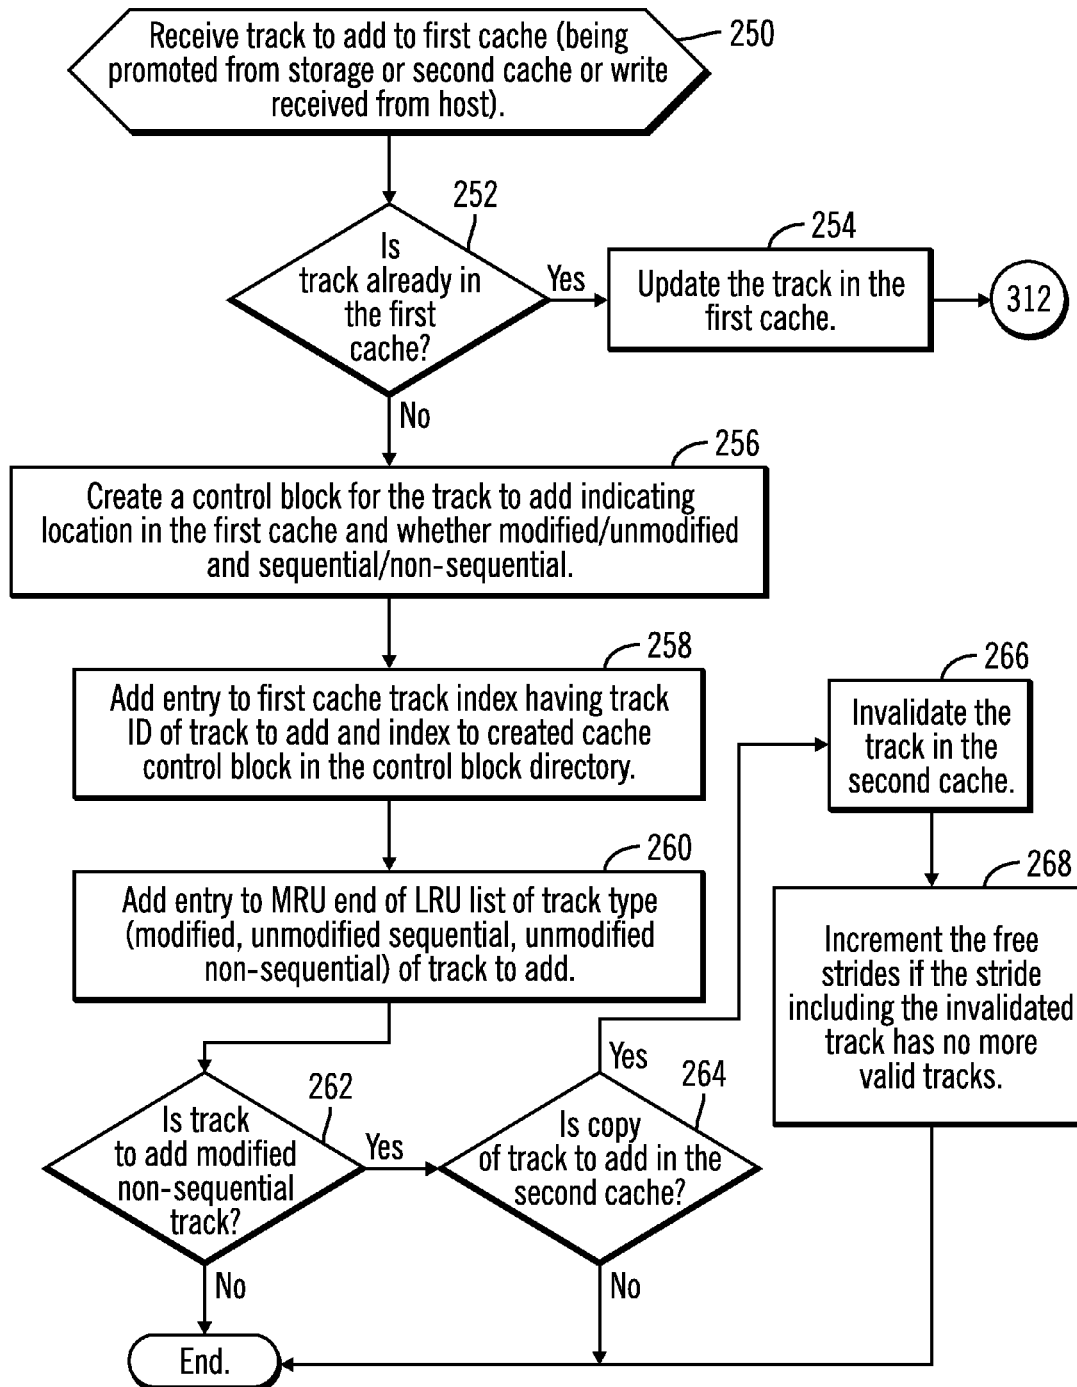


FIG. 10

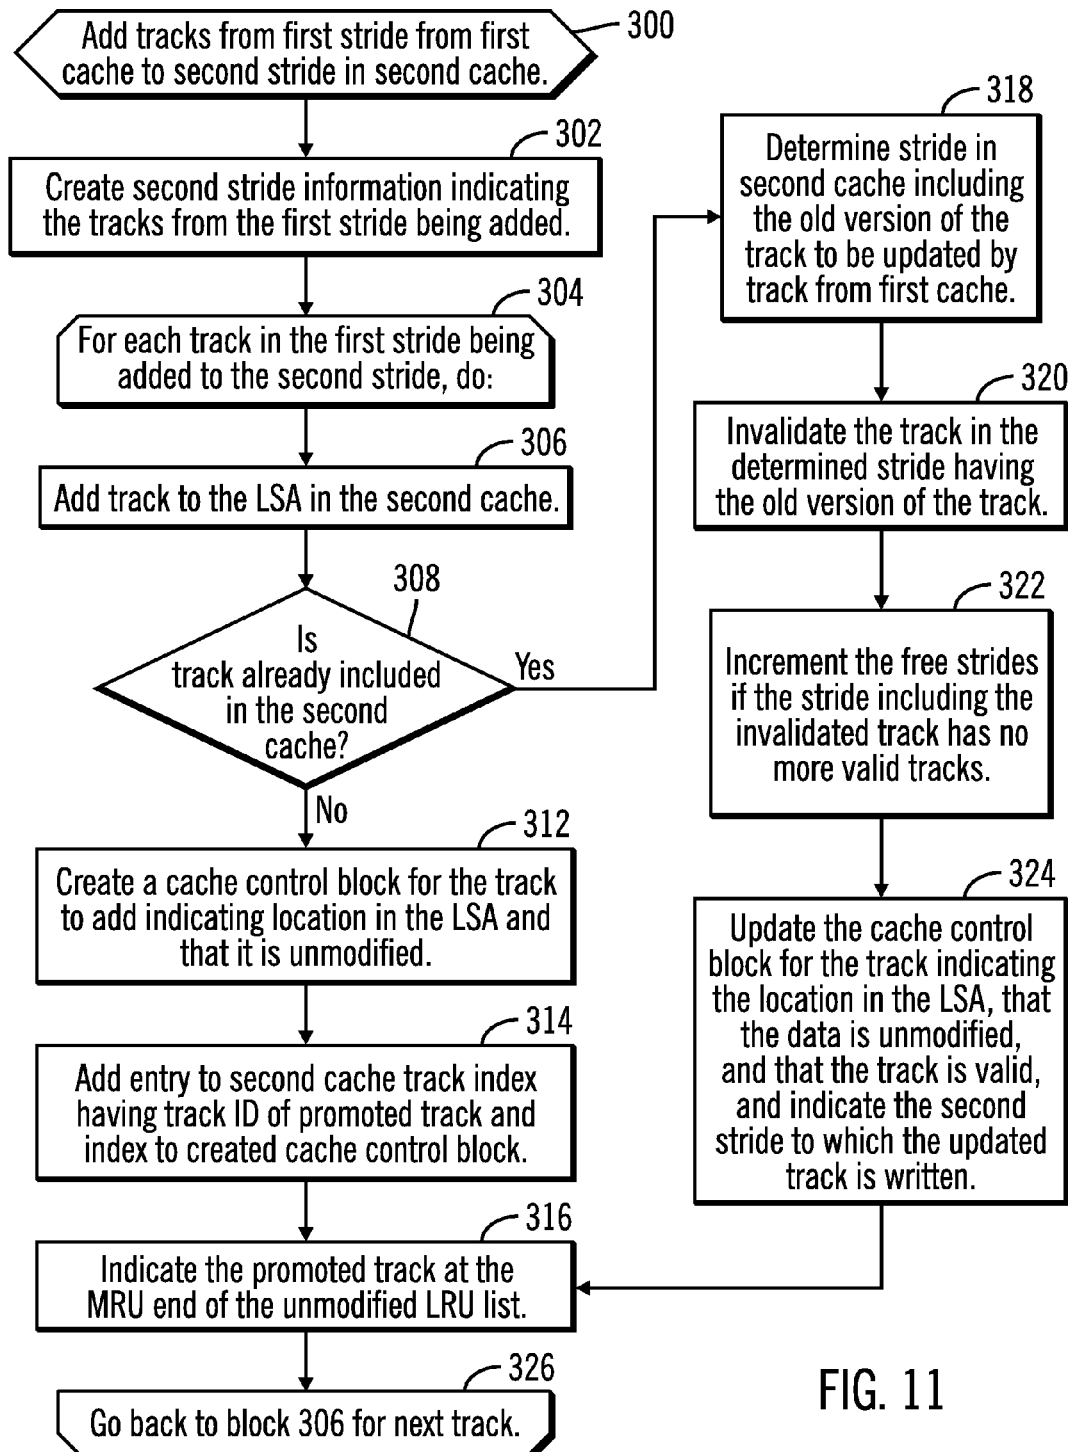


FIG. 11

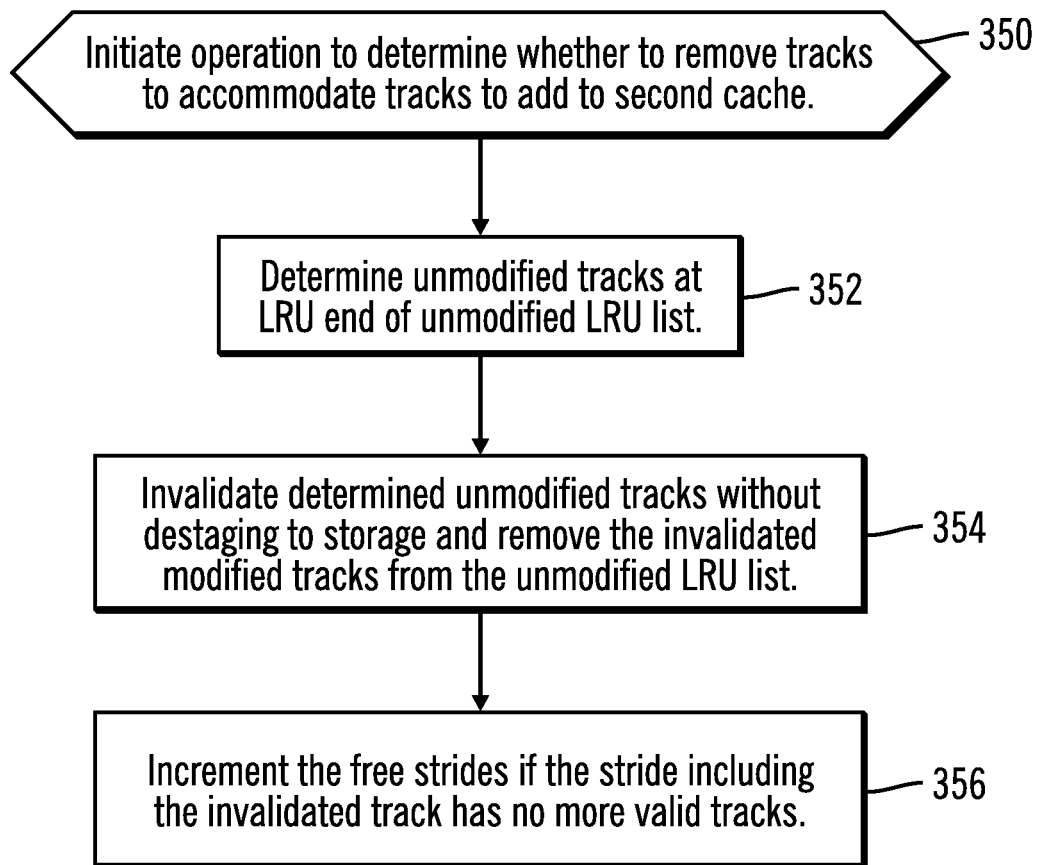


FIG. 12

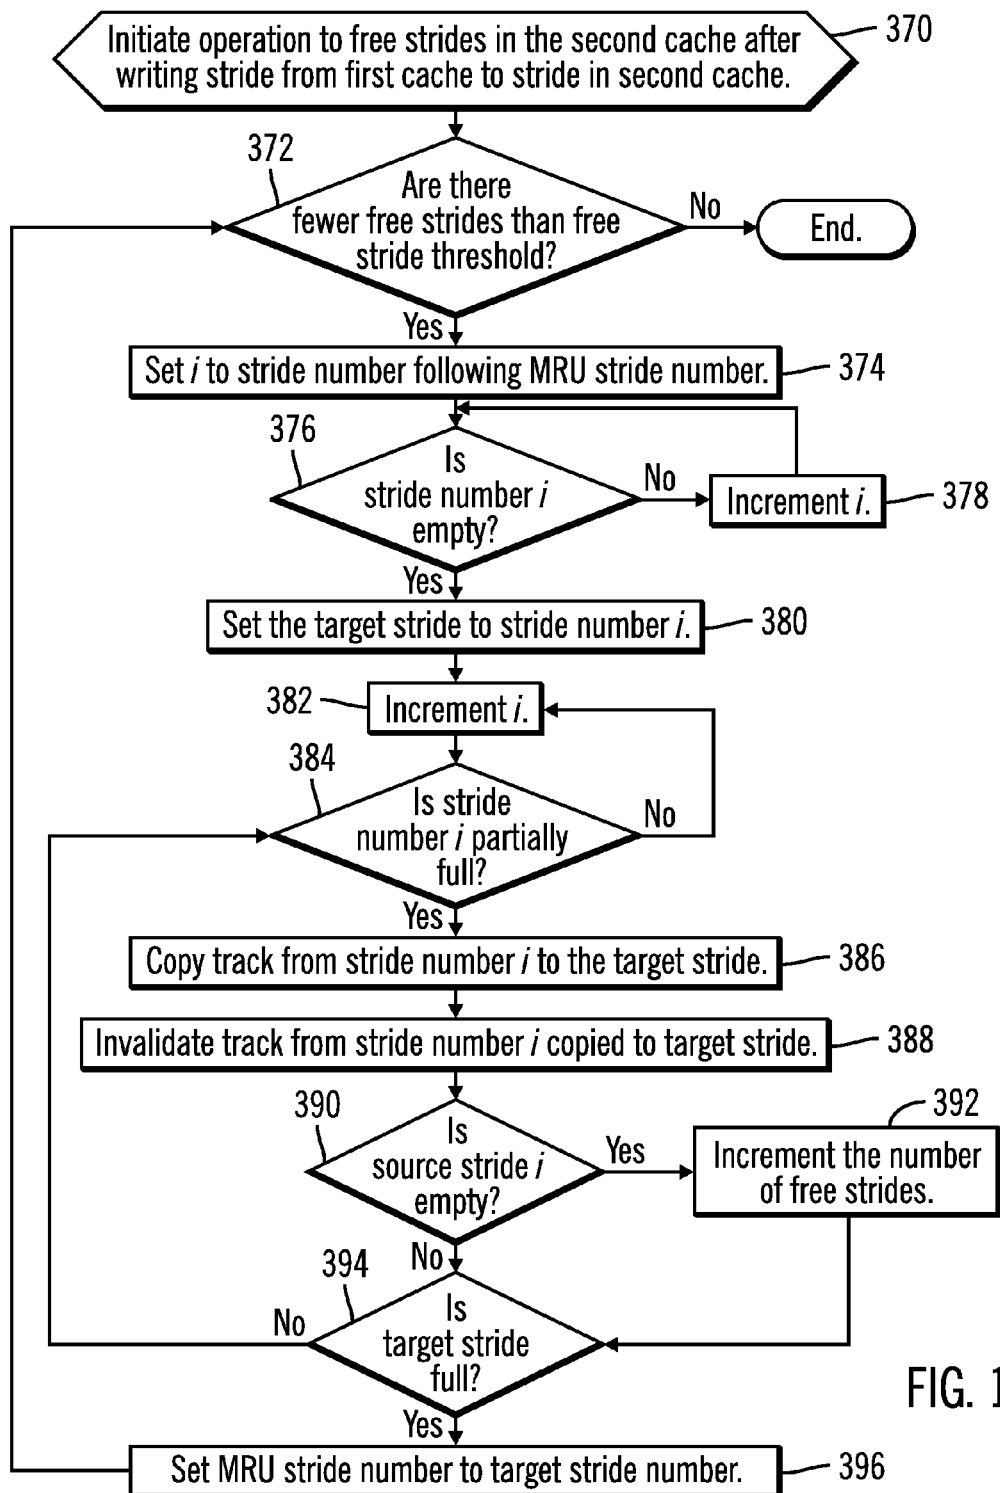


FIG. 13

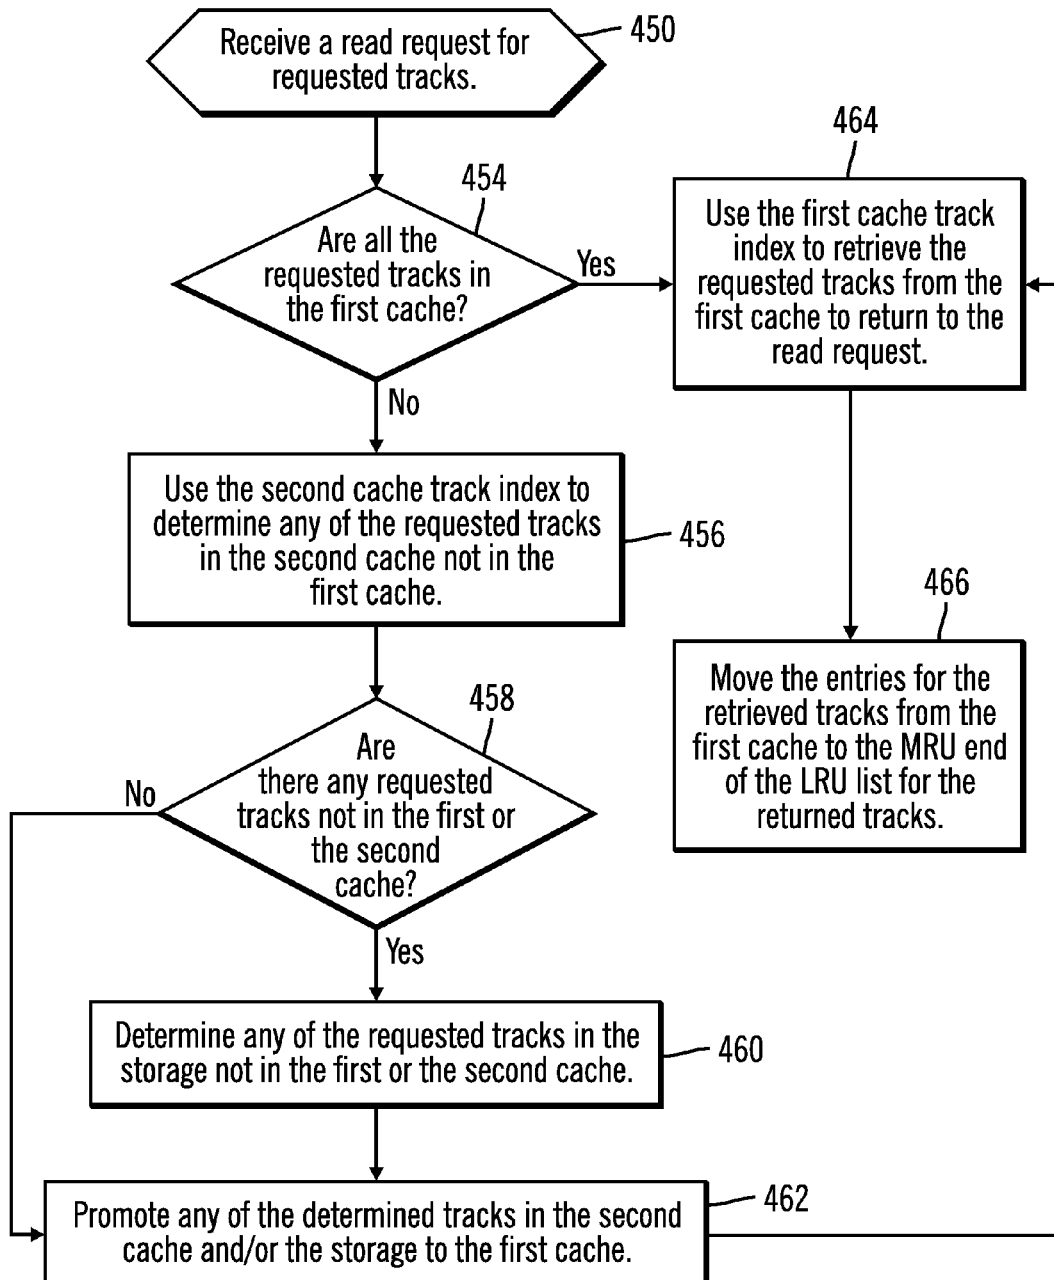


FIG. 14

1

**DEMOTING TRACKS FROM A FIRST
CACHE TO A SECOND CACHE BY USING A
STRIDE NUMBER ORDERING OF STRIDES
IN THE SECOND CACHE TO
CONSOLIDATE STRIDES IN THE SECOND
CACHE**

CROSS-REFERENCE TO RELATED
APPLICATION

This application is a continuation of U.S. patent application Ser. No. 13/779,412, filed on Feb. 27, 2013, which is a continuation of U.S. patent application Ser. No. 13/352,236, filed Jan. 17, 2012, which patent applications are incorporated herein by reference in their entirety.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to a computer program product, system, and method for demoting tracks from a first cache to a second cache by using a stride number ordering of strides in the second cache to consolidate strides in the second cache.

2. Description of the Related Art

A cache management system buffers tracks in a storage device recently accessed as a result of read and write operations in a faster access storage device, such as memory, than the storage device storing the requested tracks. Subsequent read requests to tracks in the faster access cache memory are returned at a faster rate than returning the requested tracks from the slower access storage, thus reducing read latency. The cache management system may also return complete to a write request when the modified track directed to the storage device is written to the cache memory and before the modified track is written out to the storage device, such as a hard disk drive. The write latency to the storage device is typically significantly longer than the latency to write to a cache memory. Thus, using cache also reduces write latency.

A cache management system may maintain a linked list having one entry for each track stored in the cache, which may comprise write data buffered in cache before writing to the storage device or read data. In the commonly used Least Recently Used (LRU) cache technique, if a track in the cache is accessed, i.e., a cache "hit", then the entry in the LRU list for the accessed track is moved to a Most Recently Used (MRU) end of the list. If the requested track is not in the cache, i.e., a cache miss, then the track in the cache whose entry is at the LRU end of the list may be removed (or destaged back to storage) and an entry for the track data staged into cache from the storage is added to the MRU end of the LRU list. With this LRU cache technique, tracks that are more frequently accessed are likely to remain in cache, while data less frequently accessed will more likely be removed from the LRU end of the list to make room in cache for newly accessed tracks.

There is a need in the art for improved techniques for using cache in a storage system.

SUMMARY

Provided are a computer program product, system, and method for demoting tracks from a first cache to a second cache by using a stride number ordering of strides in the second cache to consolidate strides in the second cache. Information on strides configured in the second cache

2

includes information indicating a number of valid tracks in the strides, wherein a stride has at least one of valid tracks and free tracks not including valid data. A determination is made of tracks to demote from the first cache. A first stride is formed including the determined tracks to demote. The tracks from the first stride are added to a second stride in the second cache that has no valid tracks. A target stride in the second cache is selected based on a stride most recently used to consolidate strides from at least two strides into one stride. Data from the valid tracks is copied from at least two source strides in the second cache to the target stride.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an embodiment of a computing environment.

FIG. 2 illustrates an embodiment of first cache management information.

FIG. 3 illustrates an embodiment of second cache management information.

FIG. 4 illustrates an embodiment of a first cache control block.

FIG. 5 illustrates an embodiment of a second cache control block.

FIG. 6 illustrates an embodiment of stride information.

FIG. 7 illustrates an embodiment of a second cache RAID configuration.

FIG. 8 illustrates an embodiment of a storage RAID configuration.

FIGS. 9a and 9b illustrate an embodiment of operations to demote unmodified non-sequential tracks from the first cache to promote to the second cache.

FIG. 10 illustrates an embodiment of operations to add a track to the first cache.

FIG. 11 illustrates an embodiment of operations to add tracks from the first stride to the second stride.

FIG. 12 illustrates an embodiment of operations to free space in the second cache.

FIG. 13 illustrates an embodiment of operations to free strides in the second cache.

FIG. 14 illustrates an embodiment of operations to process a request for tracks to return to a read request.

DETAILED DESCRIPTION

Described embodiments provide techniques for promoting tracks from a first cache in strides so that the tracks may be written as full stride writes to strides in the second cache to improve the efficiency of cache promotion operations. Further, while tracks are being promoted from the first cache to the second cache as strides, tracks are demoted from the second cache on a track basis according to a cache demotion algorithm, such as an LRU algorithm. To maintain free strides in the second cache to be available to store strides of tracks from the first track, strides in the second cache that are partially full, i.e., having valid and invalid tracks, may be combined into one stride. Described embodiments select source strides to merge into target strides based on a round robin technique of selecting tracks to merge from a most recently used stride number to which tracks from the first cache or second cache were written.

FIG. 1 illustrates an embodiment of a computing environment. A plurality of hosts 2a, 2b . . . 2n may submit Input/Output (I/O) requests to a storage controller 4 over a network 6 to access data at volumes 8 (e.g., Logical Unit Numbers, Logical Devices, Logical Subsystems, etc.) in a storage 10. The storage controller 4 includes a processor

3

complex 12, including one or more processors with single or multiple cores, a first cache 14 and a second cache 18. The first 14 and second 18 caches cache data transferred between the hosts 2a, 2b . . . 2n and the storage 10.

The storage controller 4 has a memory 20 that includes a storage manager 22 for managing the transfer of tracks transferred between the hosts 2a, 2b . . . 2n and the storage 10 and a cache manager 24 that manages data transferred between the hosts 2a, 2b . . . 2n and the storage 10 in the first cache 14, and the second cache 18. A track may comprise any unit of data configured in the storage 10, such as a track, Logical Block Address (LBA), etc., which is part of a larger grouping of tracks, such as a volume, logical device, etc. The cache manager 24 maintains first cache management information 26 and second cache management information 28 to manage read (unmodified) and write (modified) tracks in the first cache 14 and the second cache 18.

The storage manager 22 and cache manager 24 are shown in FIG. 1 as program code loaded into the memory 20 and executed by the processor complex 12. Alternatively, some or all of the functions may be implemented in hardware devices in the storage controller 4, such as in Application Specific Integrated Circuits (ASICs).

The second cache 18 may store tracks in a log structured array (LSA) 32, where tracks are written in a sequential order as received, thus providing a temporal ordering of the tracks written to the second cache 18. In a LSA, later versions of tracks already present in the LSA are written at the end of the LSA 32. In alternative embodiments, the second cache 18 may store data in formats other than in an LSA.

The memory 20 further includes second cache RAID configuration information 34 providing information on a RAID configuration used to determine how to form a stride of tracks to store in the second cache 18. In one embodiment, the second cache 18 may be comprised of a plurality of storage devices, such as separate solid state storage devices (SSDs), such that the strides formed of tracks from the first cache 14 are striped across the separate storage devices forming the second cache 18, such as flash memories. In a further embodiment, the second cache 18 may comprise a single storage device, such as one flash memory, such that the tracks are grouped in strides as defined by the second cache RAID configuration 34, but the tracks are written as strides to a single device, such as one flash memory, implementing the second cache 18. The tracks of strides configured for the second cache RAID configuration 34 may be written to the LSA 32 in the second cache 18 device. The second cache RAID configuration 34 may specify different RAID levels, e.g., levels 5, 10, etc.

The memory 20 further includes storage RAID configuration information 36 providing information on a RAID configuration used to determine how to write tracks from the first cache 14 or second cache 18, if the second cache 18 should store modified data, to the storage system 10, where the tracks in the destaged stride are striped across the storage devices, such as disk drives, in the storage system 10.

In one embodiment, the first cache 14 may comprise a Random Access Memory (RAM), such as a Dynamic Random Access Memory (DRAM), and the second cache 18 may comprise a flash memory, such as a solid state device, and the storage 10 is comprised of one or more sequential access storage devices, such as hard disk drives and magnetic tape. The storage 10 may comprise a single sequential access storage device or may comprise an array of storage devices, such as a Just a Bunch of Disks (JBOD), Direct Access Storage Device (DASD), Redundant Array of Inde-

4

pendent Disks (RAID) array, virtualization device, etc. In one embodiment, the first cache 14 is a faster access device than the second cache 18, and the second cache 18 is a faster access device than the storage 10. Further, the first cache 14 may have a greater cost per unit of storage than the second cache 18 and the second cache 18 may have a greater cost per unit of storage than storage devices in the storage 10.

The first cache 14 may be part of the memory 20 or implemented in a separate memory device, such as a DRAM.

The network 6 may comprise a Storage Area Network (SAN), a Local Area Network (LAN), a Wide Area Network (WAN), the Internet, and Intranet, etc.

FIG. 2 illustrates an embodiment of the first cache management information 26 including a track index 50 providing an index of tracks in the first cache 14 to control blocks in a control block directory 52; an unmodified sequential LRU list 54 providing a temporal ordering of unmodified sequential tracks in the first cache 14; a modified LRU list 56 providing a temporal ordering of modified sequential and non-sequential tracks in the first cache 14; an unmodified non-sequential LRU list 58 providing a temporal ordering of unmodified non-sequential tracks in the first cache 14; and stride information 60 providing information on strides formed of unmodified non-sequential tracks in the first cache 14 to write to the second cache 18 as a full stride write.

In certain embodiments, upon determining that the first cache 18 is full, the modified LRU list 56 is used to destage modified tracks from the first cache 14 to the storage 10 so that the copy of those destaged tracks in the first cache 18 may be discarded to make room in the first cache 18 for new modified tracks.

Once a modified non-sequential track is destaged from the first cache 14 to the storage 10, then the cache manager 24 may designate that destaged tracks as an unmodified non-sequential track in the first cache 14 and add indication of the newly designated unmodified track to the unmodified non-sequential LRU list 58, from where it is eligible to be promoted to the second cache 14. The state of the destaged modified track may be changed by updating the first cache control block 104 to indicate the destaged modified non-sequential track as unmodified in field 106. Thus, unmodified non-sequential tracks in the first cache 14 may comprise read data or modified non-sequential tracks that were destaged to the storage 10 according to the modified LRU list 56. Thus, destaged modified tracks that become unmodified tracks in the LRU list 58 may be promoted to the second cache 14 to be available for subsequent read requests. In these embodiments, the second cache 14 comprises a read only cache to cache unmodified non-sequential tracks.

FIG. 3 illustrates an embodiment of the second cache management information 28 including a track index 70 providing an index of tracks in the second cache 18 to control blocks in a control block directory 72; an unmodified list 74 providing a temporal ordering of unmodified tracks in the second cache 18; stride information 78 providing information on strides of tracks written to the second cache 18, and stride management information 80 including information for managing the strides. In certain embodiments, the strides are numbered so that stride numbers form an ordering of the strides, such as consecutive integer numbers. The stride management information may include a last stride number 80 indicating a last or highest number in the stride number ordering to assign to a stride; a most recently used (MRU) stride number 82 indicating a stride number to which tracks from the first cache 14 or second cache 18 were most recently written to fill that stride; a number of free strides 84

5

having no valid tracks; and a free stride threshold **86** indicating a minimum number of free strides to maintain. In one embodiment, the second cache **18** only stores unmodified, non-sequential tracks. In further embodiments, the second cache **18** may also store modified and/or sequential tracks.

When the second cache **18** is initialized for operations, the number of free strides **84** would be equal to the total number of strides, such as the last stride number **80** and the MRU stride number would point to the first stride number in the ordering. In this way, the stride numbers from the first stride number, e.g., 0, to the last stride number, e.g., integer **n**, provide an ordering of the strides that may be used to select strides to receive tracks from the first cache **14** or to consolidate partially filled strides to free strides.

All the LRU lists **54**, **56**, **58**, and **74** may include the track IDs of tracks in the first cache **14** and the second cache **18** ordered according to when the identified track was last accessed. The LRU lists **54**, **56**, **58**, and **74** have a most recently used (MRU) end indicating a most recently accessed track and a LRU end indicating a least recently used or accessed track. The track IDs of tracks added to the caches **14** and **18** are added to the MRU end of the LRU list and tracks demoted from the caches **14** and **18** are accessed from the LRU end. The track indexes **50** and **70** may comprise a scatter index table (SIT). Alternative type data structures may be used to provide the temporal ordering of tracks in the caches **14** and **18**.

Non-sequential tracks may comprise Online Line Transaction Processing (OLTP) tracks, which often comprise small block writes that are not fully random and have some locality of reference, i.e., have a probability of being repeatedly accessed.

FIG. 4 illustrates an embodiment of a first cache control block **100** entry in the control block directory **52**, including a control block identifier (ID) **102**, a first cache location **104** of the physical location of the track in the first cache **14**, information **106** indicating whether the track is modified or unmodified, information **108** indicating whether the track is a sequential or non-sequential access, and information **110** indicating a demote status for the track, such as no demotion, ready to demote, and demote complete.

FIG. 5 illustrates an embodiment of a second cache control block **120** entry in the second cache control block directory **72**, including a control block identifier (ID) **122**; an LSA location **124** where the track is located in the LSA **32**; modified/unmodified info **126** indicating whether the track is modified or unmodified; a valid/invalid flag **128** indicating whether the track is valid or invalid; and a stride **130** indicating the stride in which the track is included. A track in the second cache **18** is indicated as invalid if the track is updated in the first cache **14** or if the track is demoted from the second cache **18**.

FIG. 6 illustrates an instance **130** of the stride information **60**, **78** for one stride to be formed in the second cache **18**, including a stride identifier (ID) **132** and tracks **134** of the storage **10** included in the stride **132**.

FIG. 7 illustrates an embodiment of the second cache RAID configuration **34** that is maintained to determine how to form strides of tracks in the second cache **18** from the tracks in the first cache **14**. A RAID level **140** indicates the RAID configuration to use, e.g., RAID 1, RAID 5, RAID 6, RAID 10, etc., a number of data disks (**m**) **142** storing tracks of user data, and a number of parity disks (**p**) **144** storing parity calculated from the data disks **142**, where **p** can be one or more, indicating the number of disks for storing the calculated parity blocks. An unmodified parity optional flag

6

148 indicates whether parity should be calculated for unmodified non-sequential tracks in the first cache **14** being promoted to the second cache **18**. This optional flag **148** allows for only including unmodified non-sequential tracks in a stride to fill the stride with only unmodified non-sequential tracks. The stride of unmodified non-sequential tracks in the first cache **14** may be indicated in an LSA **32**, where the tracks of the stride are striped across **m** plus **p** storage devices forming the second cache **18**. Alternatively, the second cache **18** may comprise fewer than **n** devices.

FIG. 8 illustrates an embodiment of the storage RAID configuration **36** that is maintained to determine how to form strides of modified tracks in the second cache **18** to stripe across the disks of the storage **10**. A RAID level **150** indicates the RAID configuration to use, a number of data disks (**m**) **152** storing tracks of user data, and a number of parity disks (**p**) **154** storing parity calculated from the data disks **152**, where **p** can be one or more, indicating the number of disks for storing the calculated parity blocks. The stride of tracks from the second cache **18** may be striped across disks in the storage system **10**.

In one embodiment, the second cache **34** and storage **36** RAID configurations may provide different parameters or have the same parameters, such as different RAID levels, data disks, parity disks, etc.

FIGS. 9a and 9b illustrate an embodiment of operations performed by the cache manager **24** to demote unmodified non-sequential tracks from the first cache **14** to promote to the second cache **18**, where the unmodified non-sequential tracks may be selected from the LRU end of the unmodified non-sequential LRU list **58** when space is needed. Upon initiating (at block **200**) the operation to demote selected unmodified non-sequential tracks, the demote status **110** (FIG. 4) of the unmodified non-sequential tracks selected to demote is set (at block **202**) to "ready". The cache manager **24** uses (at block **204**) the second cache RAID configuration information **34** to form a first stride of tracks from the first cache **114** to promote to a stride in the second cache **18**. For instance, forming the first stride of tracks may comprise forming a stride for a RAID configuration based on a RAID configuration defined **34** for the second cache as having **n** devices including **m** devices for storing tracks of data and at least one parity device **p** to store parity data calculated from the tracks of data for the **m** devices. Further, the first stride of tracks may be striped across **n** solid state storage devices without parity to form the second stride in embodiments where the second cache comprises at least **n** solid state storage devices.

The cache manager **24** processes (at block **206**) the unmodified non-sequential LRU **58** list to determine a number of unmodified non-sequential tracks having a demote status **110** of ready in their control blocks **100**. If the cache manager **24** determines (at block **208**) that the number of unmodified non-sequential tracks is sufficient to form a stride, then the cache manager **24** populates (at block **210**) the first stride of unmodified non-sequential tracks having a demote status **110** of ready. In one embodiment, the first stride may be populated starting from the LRU end of the unmodified non-sequential LRU list **58** and use enough tracks for the data disks in stride. If (at block **212**) the RAID configuration specifies parity disks, then the cache manager **24** calculates (at block **212**) parity for the unmodified non-sequential tracks included in the stride and includes parity data (for the **p** parity disks) in the stride. If (at block **208**) there are not sufficient unmodified non-sequential tracks in the first cache **14** to fill the first stride, then control ends until there are a sufficient number of unmodified

7

non-sequential tracks having the demote ready status available to populate the first stride.

After populating the first stride (at blocks 210 and 212), control proceeds to block 220 in FIG. 9b to determine a free second stride in the second cache 18 in which to include the tracks from the first stride.

If (at block 220) the MRU stride number 82, last filled stride number in the second cache 18, is the last stride number 80 in the stride number ordering, then i, the stride number to consider, is set to the first stride number in the stride number ordering. If (at block 220) the MRU stride number 82 is not the last stride number 80, then the cache manager 24 sets (at block 224) i to the next consecutive stride number from the MRU stride number 82. If (at block 226) stride i does not have all free tracks, i.e., no valid tracks, then it cannot be used for a full stride write from the first stride of tracks from the first cache 14 and control proceeds to select the next stride number in the number ordering. To select the next stride number in the ordering, if (at block 228) stride i is the last stride number 80, then control proceeds to block 222 to consider the first stride number in the ordering. Otherwise, if stride i is not the last stride number 80, then i is incremented (at block 230) to the next stride number in the ordering.

After determining a next stride number i to consider from blocks 222, 224 or 230, the cache manager 24 determines (at block 226) whether stride i has all free tracks, i.e., no valid tracks. If not, then control proceeds to block 228 to select the next stride number in the ordering to consider. Otherwise, if (at block 226) stride i does have all free tracks, i.e., no valid tracks, then the cache manager 24 selects (at block 232) stride i as the second stride for the full stride write of tracks from the first cache 14 and writes (at block 234) the first stride of tracks from the first cache 14 to the second stride i in the second cache 18. The cache manager 24 updates (at block 236) the MRU stride number 82 to the stride number of the second stride i to which the tracks were written from the first stride at block 234. The cache manager 24 further updates (at block 238) the demote status 110 for the unmodified non-sequential tracks included in the stride as demote "complete". The cache manager 24 decrements (at block 240) the number of free strides 84 because the previously free second stride i is now no longer empty.

The described operations of FIGS. 9a and 9b utilize a round robin algorithm to select the stride in the second cache 18 to receive the full stride write from the first cache 14 by considering strides in a stride number ordering from the most recently used stride 82 previously filled.

Although the operations of FIGS. 9a and 9b are described as demoting unmodified non-sequential tracks from the first cache 14 to promote to the second cache 18, in alternative embodiments, the operations may apply to demoting different types of tracks, such as modified, sequential, etc.

With the described embodiments, the unmodified tracks from the first cache 14 are gathered and written as a stride to the second cache 18 so that one Input/Output (I/O) operation is used to transfer multiple tracks.

FIG. 10 illustrates an embodiment of operations performed by the cache manager 24 to add, i.e., promote, a track to the first cache 14, which track may comprise a write or modified track from a host 2a, 2b . . . 2n, a non-sequential track in the second cache 18 that is subject to a read request and as a result moved to the first cache 14, or read requested data not found in either cache 14 or 18 and retrieved from the storage 10. Upon receiving (at block 250) the track to add to the first cache 14, if (at block 252) a copy of the track is already included in the first cache 14, i.e., the received

8

track is a write, then the cache manager 24 updates (at block 254) the track in the first cache 14. If (at block 252) the track is not already in the cache, then the cache manager 24 creates (at block 256) a control block 100 (FIG. 4) for the track to add indicating the location 104 in the first cache 14 and whether the track is modified/unmodified 106 and sequential/non-sequential 108. This control block 100 is added to the control block directory 52 of the first cache 14. The cache manager 24 adds (at block 258) an entry to the first cache track index 50 having the track ID of track to add and an index to the created cache control block 100 in the control block directory 52. An entry is added (at block 260) to the MRU end of the LRU list 54, 56 or 58 of the track type of the track to add. If (at block 262) the track to add is a modified non-sequential track and if (at block 264) a copy of the track to add is in the second cache 18 (i.e., an older version of the track), as determined from the second cache track index 70, then the copy of the track in the second cache 18 is invalidated (at block 266), such as by setting the valid/invalid flag 128 in the cache control block 120 for the older version of the track in the second cache 18 to invalid. The cache manager 24 may further increment (at block 268) the number of free strides 86 if the stride including the track invalidated at block 266 has no more valid tracks. If (at block 262) the track to add is unmodified sequential, the second cache 18 does not include a copy (older version) of the track being added to the second cache 18 (no branch of block 264) or after incrementing the number of free strides 84 (at block 268) control ends.

The cache manager 24 can determine whether a stride 130 has no more valid tracks, i.e., is empty or free, by processing the second cache control block 120 (FIG. 5) for each track 134 in the stride.

FIG. 11 illustrates an embodiment of operations performed by the cache manager 24 to add tracks from the first stride from the first cache 14 to the second stride in the second cache 18, such as to perform the operation at block 234 in FIG. 9b. The cache manager 24 creates (at block 302) stride information 130 (FIG. 6) for the second stride indicating the tracks 134 from the first stride being added. For each track in the first stride being added, a loop of operations is performed at blocks 304 through 318. The cache manager 24 adds (at block 306) indication, such as the track ID, of the track being promoted to the LSA 32 in the second cache 18. If (at block 308) the track being added is already in the second cache 18 (i.e., an older version of the track), then the cache manager 24 determines (at block 318) the stride in the second cache 18 including the old version of the track to be updated by the track from first cache 14 and invalidates (at block 320) the track in the determined stride having the older version of the track. The cache manager 24 increments (at block 322) the number of free strides 84 if the stride including the invalidated track, invalidated at block 320, has no more valid tracks. The cache manager 24 updates (at block 324) the cache control block 120 for the track indicating the location 124 in the LSA 32, that the data is unmodified 126, and that and that the track is valid 128.

If (at block 308) the track is not already in the second cache 18, then the cache manager 24 creates (at block 312) a control block 120 (FIG. 5) for the track to add indicating the track location 124 in the LSA 32 and whether the track is modified/unmodified 126. An entry is added (at block 314) to the second cache track index 70 having the track ID of the promoted track and an index to the created cache control block 120 in the control block directory 72 for the second cache 18. From block 324 or 314, the cache manager 24

indicates (at block 316) the promoted track at the MRU end of the unmodified LRU list 74, such as by adding the track ID to the MRU end.

FIG. 12 illustrates an embodiment of operations performed by the cache manager 24 to free space in the second cache 18 for new tracks to add to the second cache 18, i.e., tracks being demoted from the first cache 14. Upon initiating this operation (at block 350), the cache manager 24 determines (at block 352) unmodified tracks in the second cache 18 from the LRU end of the unmodified LRU list 74 and invalidates (at block 354) the determined unmodified tracks without destaging the invalidated unmodified tracks to the storage 10, and also removes the invalidated unmodified tracks from the unmodified LRU list 74 and indicates the track as invalid 128 in the cache control block 120 for the track. The cache manager 24 may increment (at block 356) the number of free strides 84 if the stride including the invalidated track, invalidated at block 354, has no more valid tracks.

The unmodified tracks in the second cache 18 may comprise read tracks added to the first cache 14 or modified tracks destaged from the first cache 14. Further, the tracks selected by the cache manager 24 for demotion from the second cache 18 may be from different strides formed in the second cache 18. Further, strides in the second cache 18 may include both valid and invalid tracks, where tracks are invalidated by demoting from the second cache 18 or by the track being updated in the first cache 18.

In certain embodiments, the cache manager 24 uses different track demotion algorithms to determine tracks to demote from the first cache 14 and the second cache 18 by using separate LRU lists 58 and 74 for the first 14 and second 18 caches 18, respectively, to determine the tracks to demote. The algorithms used to select tracks for demotion in the first 14 and second 18 caches may consider characteristics of the tracks in the first 14 and second 18 caches to determine tracks to demote first.

FIG. 13 illustrates an embodiment of operations performed by the cache manager 24 to free strides in the second cache 18 to make available for strides of tracks in the first cache 14 after writing the first stride of tracks from the first cache 14 to the second stride in the second cache 18. The operations of FIG. 13 utilize a round robin algorithm to select strides to consolidate to free strides based on an ordering of the stride numbers. Upon initiating (at block 370) an operation to free strides in the second cache 18, the cache manager determines (at block 372) if the number of free strides 84 is less than the free stride threshold 86. For instance, the cache manager 24 may ensure that there are always at least two or some other number of free strides to be available for strides formed from the first cache 14 tracks. If the number of free strides is not below the threshold, then control ends. Otherwise, if (at block 372) the number of free strides is less than the threshold, then the cache manager 24 sets (at block 374) i, indicating a stride number to consider, to the stride number following the MRU stride number 82, last filled stride, in the stride number ordering.

If (at block 376) stride number i is not empty, i.e., has one or more valid tracks then it is not eligible to be the empty target stride into which strides are consolidated, and the cache manager 24 increments (at block 378) i to the next stride number in the stride number ordering. If the stride number i being incremented is the last stride number 80 in the ordering, then the increment operation at blocks 374, 378, and 382 and elsewhere sets the stride number i to the first stride number in the ordering, otherwise the stride number is incremented to the next non-last stride number. If

(at block 376) the stride number i is empty, then the cache manager 24 sets (at block 380) the target stride to stride number i into which tracks from source strides are consolidated to free the source strides. After setting the target stride, control proceeds to blocks 382 et seq. to select the source strides from which tracks are copied to the target stride, where the considered source strides to consolidate are those immediately following the target stride in the stride number ordering.

At block 382, the cache manager 24 increments i to the next stride number in the ordering. If (at block 384) the stride number i is not partially full with valid tracks, i.e., empty (having no valid tracks) or full (having only valid tracks), then it is not a suitable stride to consolidate with other strides, and the cache manager 24 proceeds back to block 382 to increment the stride number i to the next stride number in the ordering to consider for a source stride. Otherwise, if (at block 384) the stride number i is partially full (not empty nor full), then it is a suitable stride to consolidate with other strides, and the cache manager 24 copies (at block 386) a valid track from the stride number i to the target stride and invalidates (at block 388) the track from stride number i copied at block 386. After the copying, if (at block 390) source stride i is empty, i.e., having no more valid tracks to copy, then the cache manager 24 increments (at block 392) the number of free strides 84. From block 392 or if (at block 390) the source stride i being considered is not empty, then if (at block 394) the target stride is not full, control proceeds back to block 384 to copy more tracks from stride i or the next stride in the ordering to the target stride. If (at block 394) the target stride is full 394, then the cache manager 24 sets (at block 396) the MRU stride number 82 to the target stride number. From block 396, control may proceed back to block 372 to determine whether further strides need to be consolidated to provide that the number of free strides is at least the free stride threshold 86.

The described operations of FIG. 13 consider strides sequentially in an ordering for consolidation, so that after a stride is filled with tracks from the first cache 14, the cache manager 24 determines whether there a sufficient number of free strides following the stride just written in the ordering. The cache manager 24 will in a round robin manner select strides to consolidate into a stride.

FIG. 14 illustrates an embodiment of operations performed by the cache manager 24 to retrieve requested tracks for a read request from the caches 14 and 18 and storage 10. The storage manager 22 processing the read request may submit requests to the cache manager 24 for the requested tracks. Upon receiving (at block 450) the request for the tracks, the cache manager 24 uses (at block 454) the first cache track index 50 to determine whether all of the requested tracks are in the first cache 14. If (at block 454) all requested tracks are not in the first cache 14, then the cache manager 24 uses (at block 456) the second cache track index 70 to determine any of the requested tracks in the second cache 18 not in the first cache 14. If (at block 458) there are any requested tracks not found in the first 14 and second 18 caches, then the cache manager 24 determines (at block 460) any of the requested tracks in the storage 10, from the second cache track index 70, not in the first 14 and the second 18 caches. The cache manager 24 then promotes (at block 462) any of the determined tracks in the second cache 18 and the storage 10 to the first cache 14. The cache manager 24 uses (at block 464) the first cache track index 50 to retrieve the requested tracks from the first cache 14 to return to the read request. The entries for the retrieved tracks are moved (at

11

block 466) to the MRU end of the LRU list 54, 56, 58 including entries for the retrieved tracks.

With the operations of FIG. 14, the cache manager 24 retrieves requested tracks from a highest level cache 14, then second cache 18 first before going to the storage 10, because the caches 14 and 18 would have the most recent modified version of a requested track. The most recent version is first found in the first cache 14, then the second cache 18 if not in the first cache 14 and then the storage 10 if not in either cache 14, 18.

Described embodiments provide techniques to group tracks in a first cache in strides defined according to a RAID configuration for the second cache, so that tracks in the first cache can be grouped in strides to a second cache. The tracks cached in the second cache may then be grouped into strides, defined according to a RAID configuration for the storage, and then written to the storage system. The described embodiments allow full stride writes to be used to promote demoted tracks in the first cache to the second cache.

Described embodiments provide techniques to free strides of tracks in the second cache 18 by consolidating tracks from multiple source strides into a target stride based on an ordering of tracks according to stride numbers. Strides are freed in the second cache to make available for promoting tracks from the first cache 14 in strides so that the tracks may be written as full stride writes to strides in the second cache 18 to improve the efficiency of cache promotion operations. The described embodiments allow full stride writes to be used to promote demoted tracks in the first cache to the second cache in order to conserve resources by promoting an entire stride to the second cache as a single I/O operation.

Further, while tracks are being promoted from the first cache 14 to the second cache 18 as strides, tracks are demoted from the second cache 18 on a track-by-track basis according to a cache demotion algorithm, such as an LRU algorithm.

The described operations may be implemented as a method, apparatus or computer program product using standard programming and/or engineering techniques to produce software, firmware, hardware, or any combination thereof. Accordingly, aspects of the embodiments may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the embodiments may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this

12

document, a computer readable storage medium may be any tangible medium that can contain or store a program for use by or in connection with an instruction execution system, apparatus, or device.

A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electromagnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention are described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide

13

processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

The terms “an embodiment”, “embodiment”, “embodiments”, “the embodiment”, “the embodiments”, “one or more embodiments”, “some embodiments”, and “one embodiment” mean “one or more (but not all) embodiments of the present invention(s)” unless expressly specified otherwise.

The terms “including”, “comprising”, “having” and variations thereof mean “including but not limited to”, unless expressly specified otherwise.

The enumerated listing of items does not imply that any or all of the items are mutually exclusive, unless expressly specified otherwise.

The terms “a”, “an” and “the” mean “one or more”, unless expressly specified otherwise.

Devices that are in communication with each other need not be in continuous communication with each other, unless expressly specified otherwise. In addition, devices that are in communication with each other may communicate directly or indirectly through one or more intermediaries.

A description of an embodiment with several components in communication with each other does not imply that all such components are required. On the contrary a variety of optional components are described to illustrate the wide variety of possible embodiments of the present invention.

Further, although process steps, method steps, algorithms or the like may be described in a sequential order, such processes, methods and algorithms may be configured to work in alternate orders. In other words, any sequence or order of steps that may be described does not necessarily indicate a requirement that the steps be performed in that order. The steps of processes described herein may be performed in any order practical. Further, some steps may be performed simultaneously.

When a single device or article is described herein, it will be readily apparent that more than one device/article (whether or not they cooperate) may be used in place of a single device/article. Similarly, where more than one device or article is described herein (whether or not they cooperate), it will be readily apparent that a single device/article may be used in place of the more than one device or article or a different number of devices/articles may be used instead of the shown number of devices or programs. The functionality and/or the features of a device may be alternatively embodied by one or more other devices which are not explicitly described as having such functionality/features. Thus, other embodiments of the present invention need not include the device itself.

The illustrated operations of the figures show certain events occurring in a certain order. In alternative embodiments, certain operations may be performed in a different order, modified or removed. Moreover, steps may be added to the above described logic and still conform to the described embodiments. Further, operations described herein may occur sequentially or certain operations may be processed in parallel. Yet further, operations may be performed by a single processing unit or by distributed processing units.

The foregoing description of various embodiments of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be limited not by this detailed description, but rather by the claims appended hereto. The above specification, examples

14

and data provide a complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the spirit and scope of the invention, the invention resides in the claims herein after appended.

What is claimed is:

1. A method for managing data in a computer readable cache system comprising a first cache, a second cache, and a storage system comprised of storage devices, comprising: maintaining information on strides configured in the second cache, wherein a stride has at least one of valid tracks and free tracks not including valid data, wherein the strides comprise data strides populated with tracks of data, wherein the strides configured in the second cache are comprised of tracks grouped as defined by a configuration of the storage, wherein the strides of tracks formed in the second cache are striped in the storage system, and wherein the second cache caches tracks stored in the storage system;

determining tracks to demote from the first cache; forming a first stride including the determined tracks to demote;

adding the tracks from the first stride to a second stride in the second cache that has no valid tracks;

selecting a target stride in the second cache, based on a stride most recently used, to consolidate strides from at least two source strides in the second cache into the target stride; and

copying data from the valid tracks from the at least two source strides in the second cache to the target stride.

2. The method of claim 1, further comprising; maintaining indication of a number of free strides having no valid tracks;

determining whether the number of free strides is below a free stride threshold, wherein the selecting the target stride and copying the data from the valid tracks from the at least two source strides is performed in response to determining that the number of free strides is below the free stride threshold; and

invalidating the tracks in the at least two source strides copied to the target stride.

3. The method of claim 1, wherein each stride is assigned a stride number that provides an ordering of the strides, wherein the at least two source strides have stride numbers immediately following the stride number of the target stride in the ordering and wherein the at least two source strides have sufficient valid tracks to fill the target stride.

4. The method of claim 1, wherein each of the strides is assigned a stride number that provides an ordering of the strides, further comprising:

selecting for the second stride an empty stride having a next closest stride number in the ordering to the stride to which tracks from the first cache were most recently written.

5. The method of claim 1, further comprising: determining one of the tracks in one of the strides in the second cache to demote from the second cache; demoting the determined track to demote from the second cache; invalidating the determined track to demote in the second cache; and incrementing a number of free strides in response to determining that the stride including the invalidated track has no valid tracks.

6. A computer program product for managing data in a cache system comprising a first cache, a second cache having multiple storage devices for storing tracks of data, and a storage system comprised of storage devices, the computer program product comprising a non-transitory

15

computer readable storage medium having computer readable program code embodied therein that executes to perform operations, the operations comprising:

- maintaining information on strides configured in the second cache, wherein a stride has at least one of valid tracks and free tracks not including valid data, wherein the strides configured in the second cache are comprised of tracks grouped as defined by a configuration of the storage system, wherein the strides comprise data strides populated with tracks of data, and wherein the strides of tracks formed in the second cache are striped in the storage system, and wherein the second cache caches tracks stored in the storage system;
 - determining tracks to demote from the first cache;
 - forming a first stride including the determined tracks to demote;
 - adding the tracks from the first stride to a second stride in the second cache that has no valid tracks;
 - selecting a target stride in the second cache, based on a stride most recently used, to consolidate strides from at least two source strides in the second cache into the target stride; and
 - copying data from the valid tracks from the at least two source strides in the second cache to the target stride.
7. The computer program product of claim 6, wherein the operations further comprise:
- maintaining indication of a number of free strides having no valid tracks;
 - determining whether the number of free strides is below a free stride threshold, wherein the selecting the target stride and copying the data from the valid tracks from the at least two source strides is performed in response to determining that the number of free strides is below the free stride threshold; and
 - invalidating the tracks in the at least two source strides copied to the target stride.
8. The computer program product of claim 6, wherein selecting the target stride comprises selecting one of the strides in the second cache having all free tracks to be the target stride.
9. The computer program product of claim 6, wherein each stride is assigned a stride number that provides an ordering of the strides, wherein the at least two source strides have stride numbers immediately following the stride number of the target stride in the ordering and wherein the at least two source strides have sufficient valid tracks to fill the target stride.
10. The computer program product of claim 6, wherein each of the strides is assigned a stride number that provides an ordering of the strides, wherein copying valid tracks from the at least two source strides comprises:
- copying valid tracks to the target stride from strides starting from the stride having a stride number immediately following a target stride number and proceeding sequentially through each stride having a next stride number in the ordering until the target stride is full of valid tracks.
11. The computer program product of claim 6, wherein each of the strides is assigned a stride number that provides an ordering of the strides, wherein the operations further comprise:
- selecting for the second stride an empty stride having a next closest stride number in the ordering to the stride to which tracks from the first cache were most recently written.
12. The computer program product of claim 11, wherein the operations further comprise:

16

selecting for the target stride an empty stride having a next closest stride number in the ordering to the stride to which tracks from the first cache were most recently written.

13. The computer program product of claim 12, wherein the copying of the valid tracks from the at least two source strides comprises:

- copying valid tracks from a first source stride having valid tracks and a next closest stride number in the ordering to a target stride number; and continuing to copy valid tracks from source strides having valid tracks and next closest stride numbers in the ordering to the first source stride until the target stride is full.

14. The computer program product of claim 11, wherein the stride numbers in the ordering comprise sequential numbers, and wherein upon reaching a last stride number in the ordering, a next stride number to consider comprises a first stride number in the ordering.

15. The computer program product of claim 6, wherein the operations further comprise:

- determining one of the tracks in one of the strides in the second cache to demote from the second cache;
- demoting the determined track to demote from the second cache;
- invalidating the determined track to demote in the second cache; and
- incrementing a number of free strides in response to determining that the stride including the invalidated track has no valid tracks.

16. The computer program product of claim 6, wherein the first cache is a faster access device than the second cache and wherein the second cache is a faster access device than the storage devices.

17. The computer program product of claim 6, wherein the first cache comprises a Dynamic Random Access Memory (RAM), the second cache comprises a plurality of flash devices, and the storage system is comprised of a plurality of slower access devices than the flash devices.

18. A system in communication with a storage system, comprising:

- a processor;
- a first cache accessible to the processor;
- a second cache having multiple storage devices for storing tracks of data accessible to the processor;
- a non-transitory computer readable storage medium having computer readable program code embodied therein executed by the processor to perform operations, the operations comprising:

- maintaining information on strides configured in the second cache, wherein a stride has at least one of valid tracks and free tracks not including valid data, wherein the strides comprise data strides populated with tracks of data, wherein the strides configured in the second cache are comprised of tracks grouped as defined by a configuration of the storage system, wherein the strides of tracks formed in the second cache are striped in the storage system, and wherein the second cache caches tracks stored in the storage system;

- determining tracks to demote from the first cache;
- forming a first stride including the determined tracks to demote;
- adding the tracks from the first stride to a second stride in the second cache that has no valid tracks;

17

selecting a target stride in the second cache based, on a stride most recently used, to consolidate strides from at least two source strides in the second cache into the target stride; and

copying data from the valid tracks from the at least two source strides in the second cache to the target stride.

19. The system of claim **18**, wherein the operations further comprise;

maintaining indication of a number of free strides having no valid tracks;

determining whether the number of free strides is below a free stride threshold, wherein the selecting the target stride and copying the data from the valid tracks from the at least two source strides is performed in response to determining that the number of free strides is below the free stride threshold; and

invalidating the tracks in the at least two source strides copied to the target stride.

20. The system of claim **18**, wherein each stride is assigned a stride number that provides an ordering of the strides, wherein the at least two source strides have stride numbers immediately following the stride number of the

18

target stride in the ordering and wherein the at least two source strides have sufficient valid tracks to fill the target stride.

21. The system of claim **18**, wherein each of the strides is assigned a stride number that provides an ordering of the strides, further comprising:

selecting for the second stride an empty stride having a next closest stride number in the ordering to the stride to which tracks from the first cache were most recently written.

22. The system of claim **18**, wherein the operations further comprise:

determining one of the tracks in one of the strides in the second cache to demote from the second cache;

demoting the determined track to demote from the second cache;

invalidating the determined track to demote in the second cache; and

incrementing a number of free strides in response to determining that the stride including the invalidated track has no valid tracks.

* * * * *